

Relating Downlink Products to Uplink Commands in Mars Rover Operations

Paul G. Backes, Jeffrey S. Norris, Mark Powell, and Robert Steinke
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California
Paul.G.Backes@jpl.nasa.gov

Abstract— Downlink data products in Mars rover missions need to be associated with their uplink commanded target locations and the uplink commands that produced them. A system for automatically associating downlink data products with uplink commands and targets has been developed and used in a terrestrial rover field test. The approach provides insight for the design of the system to automatically associate downlink data products with the uplink products that produced them for the 2003 Mars Exploration Rover mission.

TABLE OF CONTENTS

1. INTRODUCTION
2. MOTIVATION
3. FIDO DATA TRACKING APPROACH
4. FIDO DATA TRACKING RESULTS
5. FIDO DATA TRACKING PROBLEMS
6. MER MISSION DATA TRACKING
7. CONCLUSIONS

1. INTRODUCTION

The NASA 2003 Mars Exploration Rover (MER) mission will land two rovers on Mars in early 2004. Scientists and engineers will command each rover on a daily basis which will require a very efficient Earth-based ground operations system to generate the daily uplink commands. An important part of the ground operations system is analysis of science data that has been downlinked from Mars. Two basic questions that a scientist needs answered for each data product are “Where was this data acquired?” and “What command caused this data to be acquired?” This information is critical to enable a scientist to analyze data. Providing this information efficiently enables scientific decisions to be made quickly enough to impact the daily uplink command sequence generation.

The daily rover command sequence generation process starts with the receipt of the downlink data from Mars. Engineering and science data are processed, placed in databases, and analyzed. An engineering team determines the status of the rover. Scientists evaluate scientific data and determine what activities they want the rover to perform next. Planning meetings are held to decide what activities to include in the next rover sequence and then a sequence generation team generates the rover sequence to best satisfy the desired scientific goals. The sequence is checked to verify its safety and resource usage versus constraints and then the sequence is uplinked to the rover. Reports are generated documenting the results of the various steps of the process.

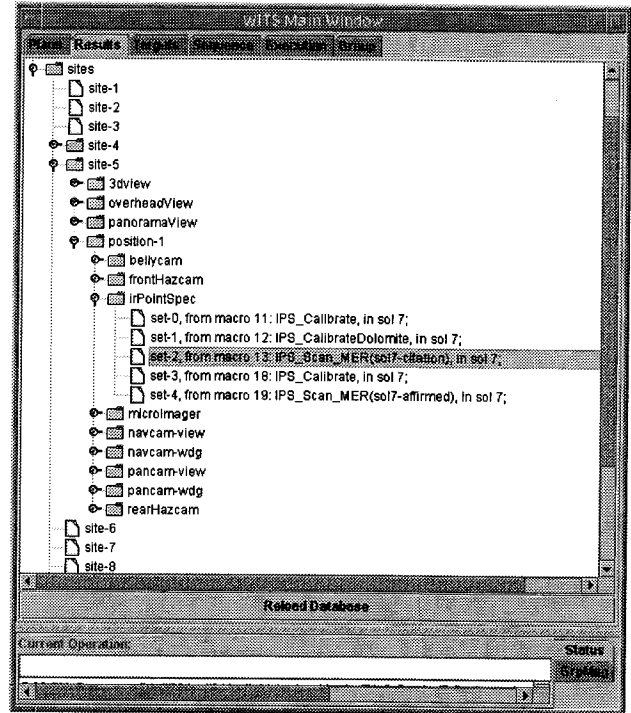


Figure 1. WITS Results Window

The Science Activity Planner (SAP) tool will be used by MER mission scientists to view downlink data and generate daily uplink science activity plans. SAP will be a version of the Web Interface for Telescience (WITS) [1]. WITS has been developed to provide Internet-based data visualization and sequence generation for Mars lander and rover missions. The round-trip data tracking capability was implemented in WITS and demonstrated during a two-week rover field test. The Field Integrated Design and Operations (FIDO) rover was at an undisclosed desert location in the California Mojave desert and commanded from JPL using WITS between April 30 and May 11, 2001. A satellite link provided communications between the primary operations center at JPL and the desert rover location. The FIDO rover is a prototype rover equipped with instrumentation designed to simulate the Athena Payload for the MER mission [2], [3]. The examples in this paper are from the FIDO field test.

Figure 1 shows the WITS Results window. The Results window provides downlink data to the user. The data is organized hierarchically by site, position, instrument, and data set. The locations that the rover traverses to are organized using sites and positions. A site is generally the area within a panorama, e.g., an area of about 20 meters radius. A position is a location that the rover has moved to within a site.




Figure 2. IPS Data Product at Citation

A user traverses the Results window tree structure to find data products of interest. When a user double-clicks on a data product name, the data product is visualized in a view specific to the data type. Figure 1 shows data set 2 for the IR Point Spectrometer instrument at position 1 of site 5 selected. The data product is shown in Figure 2.

After analyzing downlink data, scientists select targets on the terrain where they want the rover to perform science activities. They then enter the desired science activities in the sequence. Figure 3 shows WITS displays of downlink data with targets selected and a sequence being edited. The Overhead view shows imagery taken by the rover in two dimensional overhead projections; a color-coded elevation map is shown in the figure at the upper right. The Wedge view shows one image of a stereo pair. A Hazcam camera image from the front of the rover is shown in a Wedge view in the bottom right in the figure. The 3D View shows a solid model of the rover and terrain. Other views not shown in the figure are also available, e.g., a Panorama view that provides a mosaic of images taken by rover mast-mounted cameras.

A user creates a science target by selecting a pixel in an image. The associated 3D location on the terrain is then generated by WITS and the user can specify a name for this 3D location. The targets specified by the users are shown as circles with the target names next to them in the Overhead and Wedge views of Figure 3. The targets are shown as cones in the 3D view.

The sequence is built using the Sequence window shown in the upper left of Figure 3. Available commands are provided in the left column and users enter them at desired positions in the sequence on the right and specify the command parameters such as target name. One or more commands are encapsulated in macros in a sequence. Users can simulate the sequence and check rules and resource utilization versus allocations before sending the sequence to the rover for execution.

When the sequence is sent for execution on the rover, the commands are converted to the language of the rover. For the FIDO rover this means converting the higher-level WITS commands into the lower-level FIDO rover commands. For example, WITS commands can include target names as parameters but FIDO commands have only numeric parameters. WITS converts the sequence commands before sending the sequence to the FIDO rover for execution.

2. MOTIVATION

Next to the data set name in Figure 1 is the sequence name, "sol 7", and the macro ID and name that produced the data product. Also, for commands with a terrain target location to take data at, the target name is shown in parenthesis. The sequence, macro name, and target name for a data product were not available in previous versions of WITS. Without this information it was difficult for a scientist user to figure out

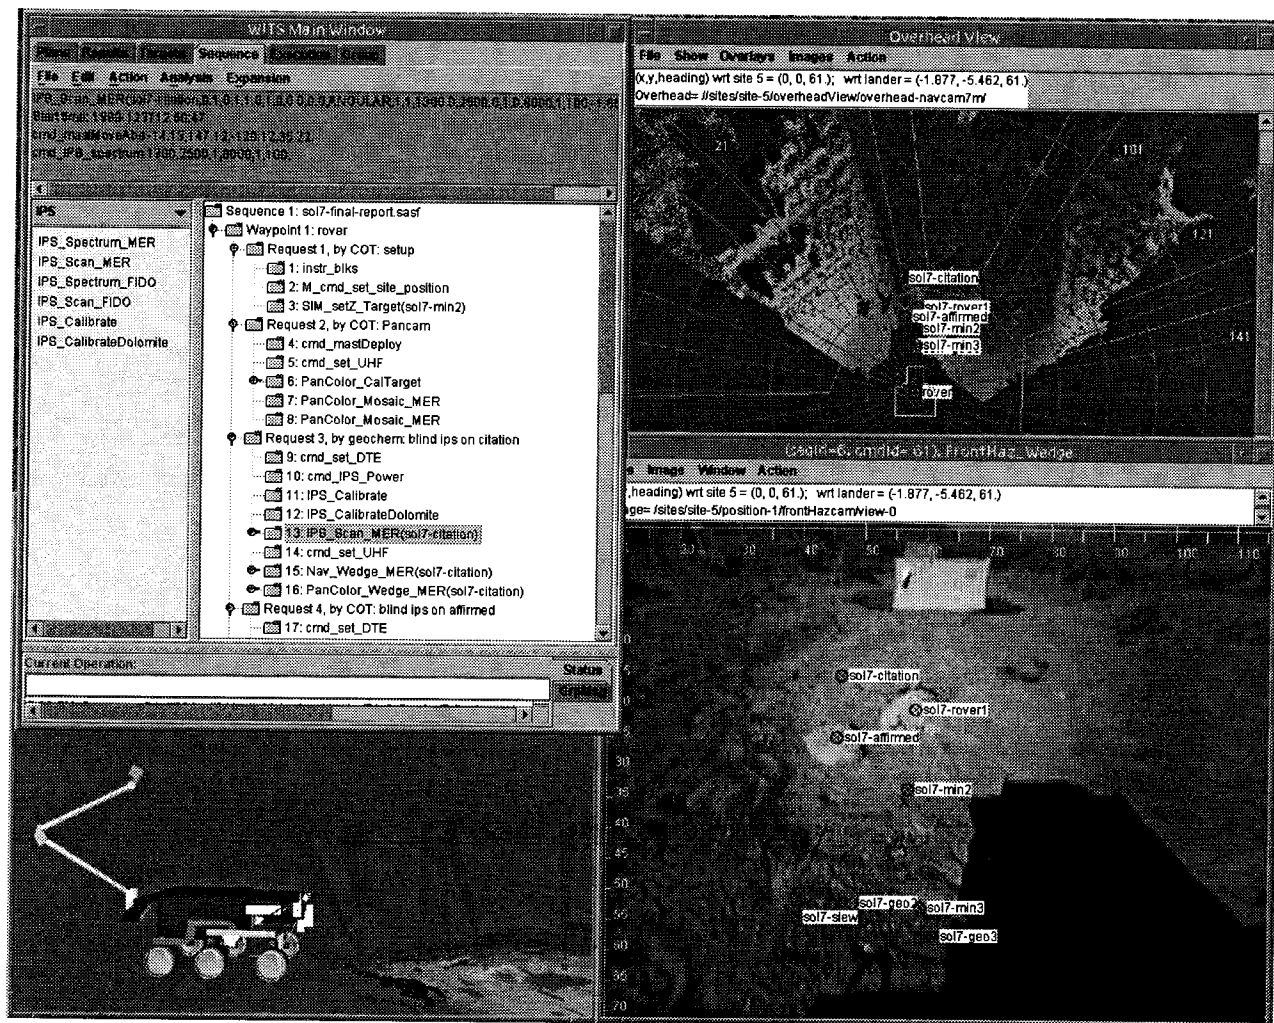


Figure 3. WITS Views and Sequence Window

where a data product was acquired and what parameters were used in the instrument command that produced the data. It was a very time consuming and frustrating process for a scientist to try to figure out this information. It was clear that a means for automatically providing this information was necessary in order to enable scientists to analyze downlink data fast enough to use their analysis results in their rover activity planning for the next day's rover activities for the MER mission.

Experience from previous FIDO rover field tests showed that relating downlink data with the uplink information that produced it is also needed in other parts of the science analysis and planning process. When looking at a data product the user needs to know its associated uplink sequence, macro, and target information. Web browser based reports will be generated during a mission to provide information on uplink sequences and downlink data. The uplink sequence information needs to be related to the downlink data products that they produced in the reports. An earlier version of automated report generation did not associate uplink information with downlink data products [4]. Associating uplink information with downlink data products has now been developed and is described below.

3. FIDO DATA TRACKING APPROACH

To enable relating uplink information with downlink data products, a sequence identifier and a command identifier were included with each data product downlinked from the FIDO rover. The sequence identifier was uplinked with the sequence. The on-board rover system incremented a command counter each time a command was executed and included the sequence identifier and current command counter with data products that a command produced. The inclusion of sequence and command identifiers with the downlink data products for FIDO rover operations is depicted in Figure 4.

The downlink data processing system put information for each data product in a separate directory in the WITS database along with an information file [5]. The information file included the sequence identifier and command identifier.

After all downlink data from the rover in one communication session was received, processed, and placed in the database, another processing step was done that generated the data to relate the uplink information with the downlink data both for use in WITS and for automatically generated HTML reports. For each data product, an uplink.info file was created in its




Figure 4. Data tracking for FIDO

database directory that contained its associated uplink information including sequence name, target name, macro, and command.

To relate uplink information with data products, a sequence is loaded and a sequence table is made that is keyed by command identifiers and for each command the sequence name, target location, macro, and command string are stored. Command identifiers are created in the same way as is done on the rover: a command counter is incremented for each command in the sequence. For each downlink data product, the command identifier is read from its downlink information file and the command identifier is used as the key in the sequence table to find the associated uplink data. This uplink data is then written out in the data product's uplink.info file in its database directory.

A downlink table is also created and keyed by command number. The paths to all data products produced by a command are stored with the command key.

For automated HTML report generation, HTML document reports are generated by stepping through a sequence and writing out sequence information and paths to data that was generated by each step of the sequence. For a sequence, the commands are incremented and path information to all data products that a command produced are read from the downlink table using the command identifier as the table key.

4. FIDO DATA TRACKING RESULTS

The technology to associate downlink data with uplink data products was used in a two week FIDO field test where FIDO was in the California Mojave desert and WITS was used at JPL to generate command sequences. A satellite link provided communications between JPL and the desert location. The approach described above was used to include the sequence identifier and command number with downlink data products and then automatically generate the data to relate downlink products with uplink information and generated information for reports.

The improvement to the Results window was described in Section 2. and shown in Figure 1. The Results window dis-

plays downlink data organized by rover site, position, instrument, and data set. When the Results window is loaded, the uplink.info file for a data product is read and the information is displayed next to the data set in the Results window. A user can immediately see what sequence, macro, and target produced the data. Prior to this development, only the name of the data product was provided and scientist users complained that they could not figure out what sequence, macro, command, and target produced the data. Another associated feature that was developed is state visualization for a data product. When a user clicks on the name of a data product in the Results window, the rover state is updated in all views of WITS with the rover position and configuration when it acquired the data, e.g., the mast, instrument arm, position and heading are all updated. A 3D view shows the rover as a solid model on a terrain, a Panorama view shows a wire frame rover in images taken by the rover, and an Overhead view shows the rover as a 2D icon in an overhead projection of image data. The user can then quickly see where the rover went and what data was produced at each rover position.

Another improvement in relating uplink data with downlink products is in the display of uplink information in the data product view. When a data product view is opened, information from the uplink.info file is provided in the view. In Figure 2, the sequence identifier, command number, macro name, and target name are shown at the top of the view. The uplink information is also shown in the text area above the plot in the view. Previously this uplink information was not provided with the data product which made it difficult for a user to figure out the context of the data.

The automatically produced HTML reports were also greatly improved. An operations report was provided for the field test that collected information about the field test. The operations report system was developed for the May 2000 FIDO field test of a year earlier [4], but it lacked relating uplink information with downlink data products. The operations report for the April-May 2001 FIDO field test is shown in Figure 5. The left column provides links to various types of information about the field test. Sequence reports give detailed information about the uplink sequences. The results link allows a user to traverse the database and look at data products. The sequence reports and results capabilities allow looking at detailed uplink and downlink information, but they do not provide a means to relate the uplink and downlink information.

The SeqPos Reports link of the operations report provides access to SeqPos reports. Selecting the SeqPos Report link causes the SeqPos reports to be listed in the middle column, as shown in Figure 5. For each uplink sequence, two SeqPos reports are generated, one at the macro level and one at the command, or step, level. The macro level SeqPos report lists data products by macro in a sequence. The step level SeqPos report lists data products by command in a sequence. There may be multiple commands in one macro, so providing the two options allows the user to decide what level of sequence detail to use.

Figure 6 is the SeqPos report that is presented when the user selects the "sol7-final.sasf (Step)" SeqPos link in the middle column of the operations report of Figure 5. The sequence el-

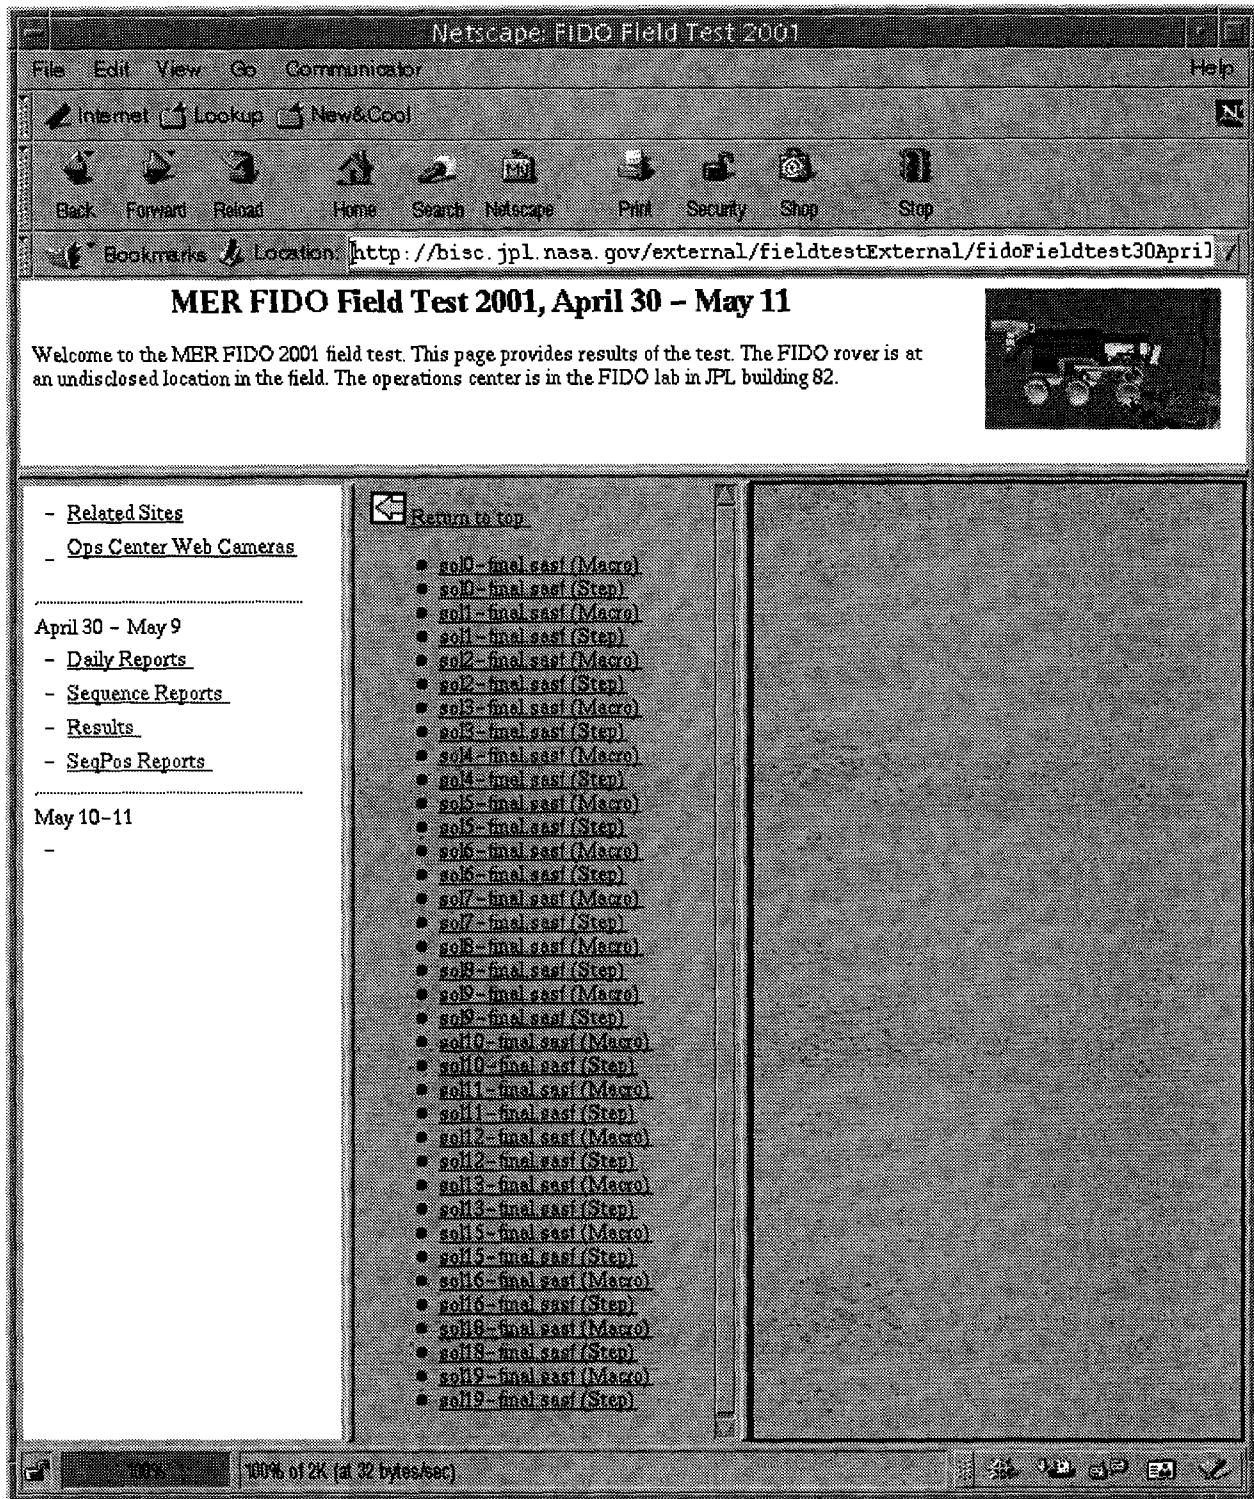


Figure 5. Operations Report

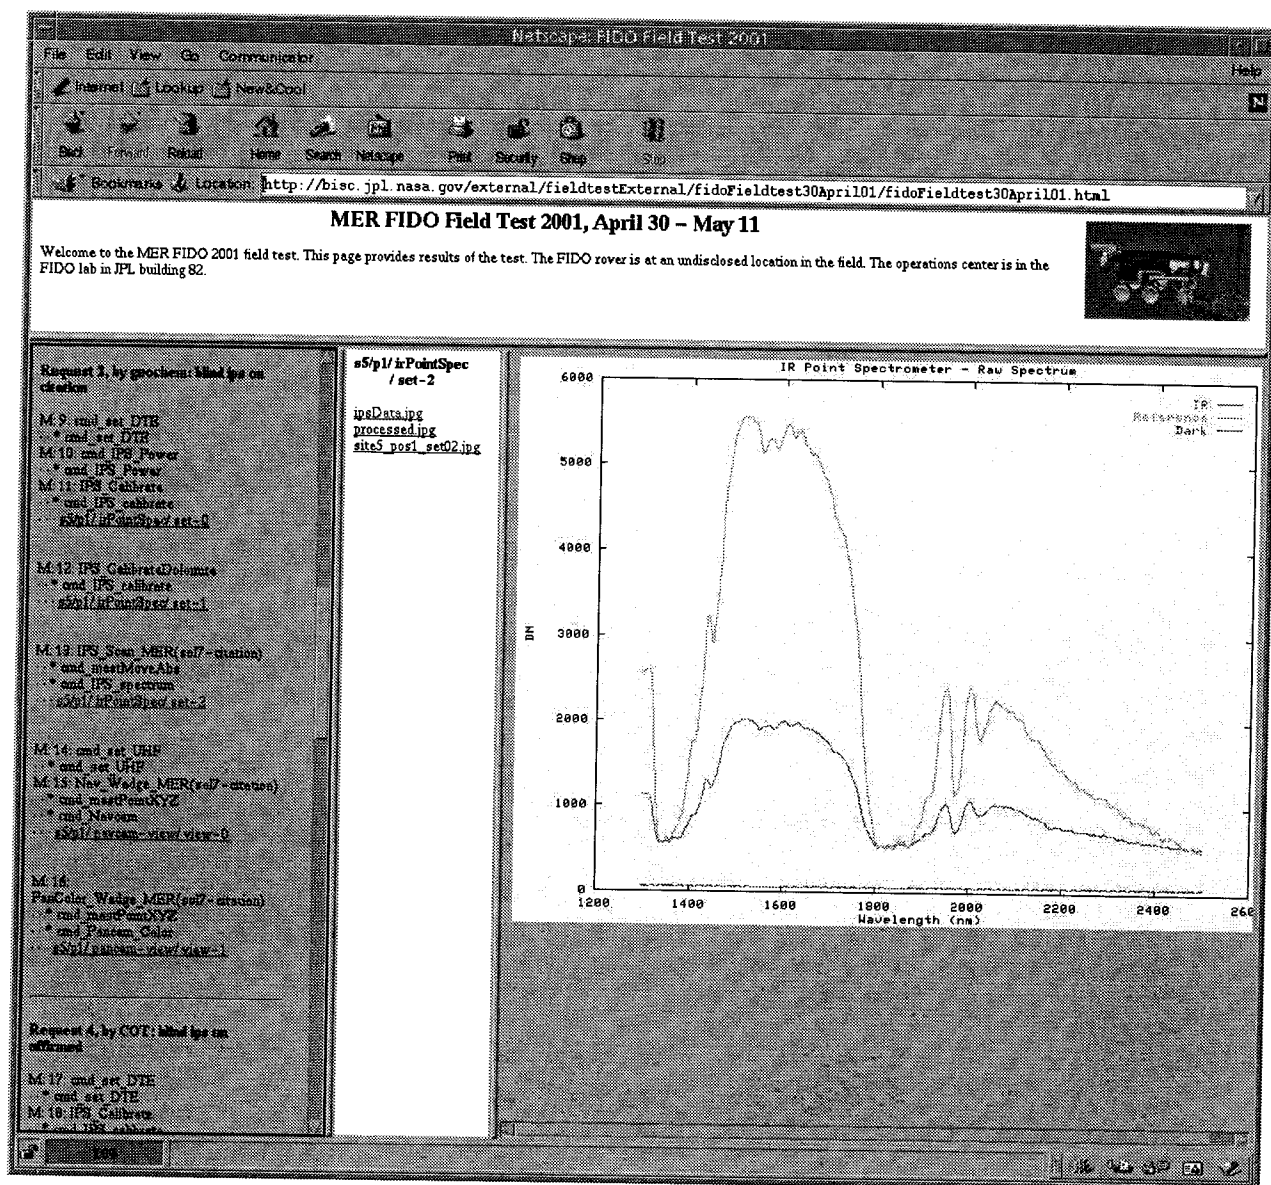


Figure 6. Sol7 SeqPos Report

ements are shown in the left column with links to the resulting data products listed under them. The figure shows the results of the sol 7 sequence that is in the Sequence window of Figure 3. The results of Request 3 and part of Request 4 of the sequence are shown. Selecting a data set in the left column causes the available products of that data set to be listed in the middle column. Selecting the data product in the middle column causes the data product to be displayed in the right column. The same IPS data product as was viewed in WITS in Figure 2 is shown in the report of Figure 6.

5. FIDO DATA TRACKING PROBLEMS

One serious problem occurred using the approach to relating uplink information to downlink data products. The FIDO rover sometimes generated incorrect command numbers. This was usually due to the rover aborting the sequence

execution and then having the sequence restarted at the point it was aborted. When the sequence was restarted, the command counter was reset, but the sequence was restarted at a command somewhere in the sequence and not at the beginning, so the command numbers sent with the downlink data were wrong. There was not enough time during the field test to manually fix the command numbers, so sequences that had this problem had incorrect uplink information stored with their data products and the SeqPos reports had incorrect associations of data products with sequence steps. It was concluded that in the future command identifiers should be uplinked with each command to prevent the problem of incorrect generation of command numbers on-board.

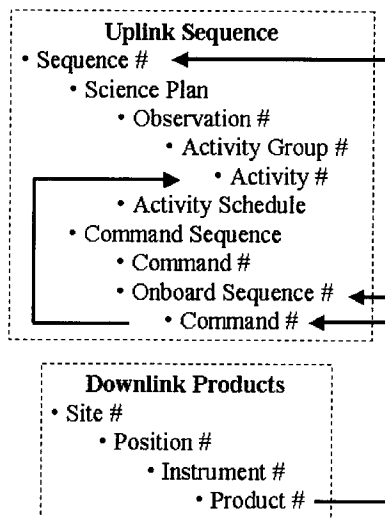


Figure 7. Data tracking for MER

6. MER MISSION DATA TRACKING

The approach for associating uplink information with downlink data products is currently being designed for the MER mission. There are a couple of differences between FIDO rover system and the MER rover system that makes the problem more complicated. First, the uplink sequence has a different, and more complicated, structure for MER, as shown in Figure 7. The sequence has two parts, the science plan and the command sequence. The science plan has observation definitions and an activity schedule. The command sequence is the sequence that will be uplinked to the rover (the science plan is not uplinked to the rover). Commands in the command sequence will have a parameter that specifies what activity they satisfy. This parameter will probably be removed from the uplink version of the command sequence. It is needed in the saved version of the command sequence in order to relate the command to its associated activity. Another version of the command sequence is produced in another format that is actually sent to the rover. That version will have the activity parameter removed.

The command sequence can utilize subsequences that have been stored on-board. A downlink data product needs to have enough information with it to enable the ground data system to automatically find the sequence, science plan, observation, activity group, activity, and command that produced it.

It is recommended that three identifiers be downlinked with each MER rover science data product: the sequence identifier, the on-board subsequence identifier, and a command identifier. The identifiers should be uplinked with the command sequence. A sequence will have a command that specifies the sequence identifier. Each command that is part of the main uplinked sequence will have an identifier uplinked with the command. If an on-board subsequence is to be used, then the command in the main sequence that calls it will give the on-board subsequence a unique identifier for each time it is run. Commands that are part of on-board subsequences will have identifiers that were specified when the on-board subsequence was uplinked. The on-board subsequence identifier value downlinked with a data product will be set to a default

value if the command is part of the main sequence. With this information, the ground data system will be able automatically associate uplink information with all downlink data products and the system features described in this paper will be available for the MER mission users.

7. CONCLUSIONS

Automatically associating uplink information with downlink data products is important for Mars rover mission operations. An approach has been developed and tested in a field test for associating uplink information with downlink data products both for use within the WITS tool and in automatically generated HTML reports. The results were very useful to science users during the field test. One important problem occurred with the approach during the field test. Automatic incrementing of an on-board command counter resulted in incorrect command numbers being downlinked with data products after the rover execution was aborted and then restarted. Due to this problem, it is recommended that command identifiers be uplinked with each command for the MER mission.

ACKNOWLEDGEMENTS

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

- [1] Paul G. Backes, Kam S. Tso, Jeffrey S. Norris, Gregory K. Tharp, Jeffrey T. Slostad, Robert G. Bonitz, and Khaled S. Ali. Internet-based operations for the mars polar lander mission. In *Proceedings IEEE International Conference on Robotics and Automation*, pages 2025–2032, San Francisco, California, April 2000.
- [2] R. Arvidson, Paul Backes, E. Baumgartner, D. Blaney, L. Dorsky, A. Haldemann, R. Lindemann, P. Schenker, and S. Squyers. Fido: Field-test rover for 2003 and 2005 mars sample return missions. In *30th Lunar and Planetary Science Conference*, Houston, Texas, March 15–19 1999.
- [3] P.S. Schenker, E.T. Baumgartner, L.I. Dorsky, P.G. Backes, H. Aghazarian, J.S. Norris, T.L. Huntsberger, Y. Cheng, A. Trebi-Ollennu, M.S. Garrett, B.A. Kennedy, and A.J. Ganino. Fido: a field integrated design & operations rover for mars surface exploration. In *Proceedings 6th International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS-'01)*, Montreal, Canada, June 18–21 2001.
- [4] Paul G. Backes and Jeffrey S. Norris. Automated rover sequence report generation. In *Proceedings IEEE Aerospace Conference*, Big Sky, Montana, March 2001.
- [5] Jeffrey S. Norris, Paul G. Backes, and Eric T. Baumgartner. Ptep: The parallel telemetry processor. In *Proceedings IEEE Aerospace Conference*, Big Sky, Montana, March 2001.




Paul G. Backes is a technical group leader in the Autonomy and Control Section at Jet Propulsion Laboratory, Pasadena, CA, where he has been since 1987. He received the BSME degree from U.C. Berkeley in 1982, and MSME in 1984 and Ph.D. in 1987 in Mechanical Engineering from Purdue University. He is currently responsible for distributed operations research for Mars lander and rover missions at JPL. Dr. Backes received the 1993 NASA Exceptional Engineering Achievement Medal for his contributions to space telerobotics (one of thirteen throughout NASA), 1993 Space Station Award of Merit, Best Paper Award at the 1994 World Automation Congress, 1995 JPL Technology and Applications Program Exceptional Service Award, 1998 JPL Award for Excellence and 1998 Sole Runner-up NASA Software of the Year Award. He has served as an Associate Editor of the IEEE Robotics and Automation Society Magazine.



Jeffrey S. Norris is a computer scientist and member of the technical staff of the Autonomy and Control Section at Jet Propulsion Laboratory. He specializes in software engineering for telerobotics, distributed operations, machine vision, and large scale interfaces. He received his Bachelor's and Master's degrees in Electrical Engineering and Computer Science from MIT. While an undergraduate, he worked at the MIT Media Laboratory on data visualization and media transport protocols. He completed his Master's thesis on face detection and recognition at the MIT Artificial Intelligence Laboratory. He now lives with his wife in Azusa, California.

Mark Powell is a member of the technical staff in the Autonomy and Control Section at the Jet Propulsion Laboratory, Pasadena, CA since 2001. He received his B.S.C.S. in 1992, M.S.C.S. in 1997, and Ph.D. in Computer Science and Engineering in 2000 from the University of South Florida, Tampa. His dissertation work was in the area of advanced illumi-



nation modeling, color and range image processing applied to robotics and medical imaging. At JPL his area of focus is science data visualization and science planning for telerobotics. He is currently serving as a software and systems engineer, contributing to the development of science planning software for the 2003 Mars Exploration Rovers and the JPL Mars Technology Program Field Integrated Design and Operations (FIDO) rover task. Dr. Powell received the 2001 Outstanding Dissertation Award from the University of South Florida. He, his wife Nina, and daughters Gwendolyn and Jacquelyn live in Pasadena, CA.

Robert Steinke is a computer scientist and member of the technical staff of the Autonomy and Control Section at the Jet Propulsion Laboratory. At JPL, his work is focused in the areas of distributed operations for Mars rovers and landers, secure data distribution, distributed data consistency, and communication support for group collaboration. He received his B.S. in 1995 and M.S. in 1997 in Computer Science from U.C. Santa Barbara, and his Ph.D. in 2001 in Computer Science from the University of Colorado at Boulder. As a graduate student he completed his Master's thesis on a serializable lazy update scheme for distributed databases, and his dissertation on distributed shared memory consistency models, and received a graduate teacher certificate from Boulder for his training and work as an instructor. He now lives in Pasadena with his wife Mollie and daughter Adelaide.